

OPTIMISTIC PROCESSING OF NETWORK FRAMES TO REDUCE LATENCY

BACKGROUND

5 1. Field of the Present Invention

The present invention relates to the field of data processing networks and more particularly to the processing of frames received by a server on the network.

2. History of Related Art

10 In the field of data processing networks, the client-server model is well known. In a client-server model, the client may comprise an application such as a conventional web browser running on any manner of computing device connected to a network such as the Internet. The client device may be a desktop or laptop personal computer, a network computer, an Internet enabled phone or personal digital assistant (PDA), and so forth. The server may comprise an application such as a web server application running on a second computing device (or set of computing devices) such as a server appliance that is also connected to the network. In a typical client server transaction or session, the client initiates a request for information that is delivered to the server over the network. As is well known, the request may travel over the network as one or more frames of information formatted according to a network protocol supported by both the client and server. Typically, the formatting of frames according to a network protocol includes appending one or more frame headers onto the data (or a portion thereof) representing the request at the client end. The number of frame headers appended onto a request may vary dependent upon the particular network protocol in use. The Open System Interconnection (OSI) Model defines seven layers that describe how applications running on network-aware devices
25 communicate with each other. A network frame may include a header corresponding to each of these layers (except perhaps the physical layer). When a frame arrives at its destination, the server device processes the header information prior to delivering the application data to the server application.

30 Among the most prevalent type of client-server communication is a Hypertext Transfer Protocol (HTTP) formatted request delivered to a server over a Transmission Control Protocol/Internet Protocol (TCP/IP) compliant network. The TCP/IP suite of protocols is the

basic communication language of many networks including the Internet. For more information regarding TCP/IP, the reader is directed to M. Murhammer et al., *TCP/IP Tutorial and Technical Overview*, available online at www.redbooks.ibm.com (#GG24-3376-05) and incorporated by reference herein. When an HTTP request is received at the server end, the lower level headers including the IP header and the TCP header are processed to verify, among other things, that the frame has arrived at the correct destination and has not been corrupted during transmission over the network. When the lower level headers have been processed, the HTTP request is passed to a web server to process the actual request. Thus, the delivery of the request to the server application is typically delayed until the network headers are processed. It would be desirable to implement a method and system that reduced or eliminated this delay in an effort to improve network performance.

SUMMARY OF THE INVENTION

The problems identified above are in large part addressed by a system and method in which the server device processes the lower level headers (referred to herein as the “network portion”) of a frame substantially in parallel with the processing of the application portion of the frame. The application portion of the frame, which may include an HTTP request is forwarded to the server application such as a web server, while the network portion of the frame is processed. If the processing of the network portion determines that the frame was mis-delivered or is corrupted, the response to the HTTP request is aborted, otherwise the response is processed and returned to the client. By optimistically assuming that the request was delivered correctly, the present invention leverages the parallel processing capabilities available on many server appliances and improve response time without incurring any substantial performance penalty.

BRIEF DESCRIPTION OF THE DRAWINGS

Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

FIG 1 is a block diagram of selected portions of a data processing network suitable for use in conjunction with the present invention;

FIG 2 is a conceptualized illustration of a network frame; and

FIG 3 is a block diagram of a system and method for processing client request in a data processing network according to one embodiment of the invention.

While the invention is susceptible to various modifications and alternative forms, specific
5 embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description presented herein are not intended to limit the invention to the particular embodiment disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

10

DETAILED DESCRIPTION OF THE INVENTION

Turning now to the drawings, FIG 1 depicts selected features of a data processing network 100 according to one embodiment of the present invention. Portions of the network and systems described herein may be implemented as a set of computer executable instructions (computer software). For such portions, the instructions may reside on any suitable computer readable memory or storage facility. During execution of the instructions, for example, the software, or portions thereof, may reside in the system memory (RAM) or cache memory of a data processing device. At other times, the software be stored on a permanent or quasi-permanent storage facility such as a hard disk, floppy diskette, VD ROM, DVD, magnetic tape or other suitable storage facility. In the embodiment depicted in FIG 1, network 100 includes a client device 102 connected to a network 104. Client device 102 may be implemented as substantially any network-aware data processing device including, as examples, desktop and laptop personal computers, workstations, network computers, internet enabled phones and
25 PDA's, etc.

Client device 102 includes a client application program that enables client device 102 to communicate with other devices connected to the network. In one embodiment, the client application program is a conventional web browser indicated in FIG 1 by reference numeral 103. Using browser 103, a user of client 102 can initiate requests for information from other network
30 devices. A common example of a request for information is an HTTP GET request that occurs when the user specifies a Universal Resource Locator (URL) through browser 103.

The GET request, or any other request for information initiated by a user of browser 103, is formatted according to the network protocol by the operating system and/or firmware of client device 102 and delivered to network 104 typically via a network interface card (not shown). Network 104 may represent a private, local area network such as an Ethernet network or a wide area network such as the Internet.

Referring to FIG 2, a conceptualized representation of a frame 200 suitable for transmission over an OSI model compliant network 104 is presented. As illustrated, frame 200 includes a data field 202 (also referred to as a payload) an application layer header 204, a transport layer header 206, and a network layer header 208. Specific implementations of network 104 may require additional headers as well. In an embodiment in which network 104 is a TCP/IP network, such as the case in which network 104 represents the Internet, the transport layer header 206 represents the TCP header, and the network layer 208 represents the IP header. If the frame is issued as part of a web based request from web browser 102, the request is likely an HTTP formatted request in which case the application layer header 204 of frame 200 is an HTTP header.

Returning now to FIG 1, a server device 110 that is the target of the request initiated by client 102 is also connected to network 104. Typically, server device 110 includes a network interface card (NIC) 112 connected to one or more central processing units (CPUs) 111 through a intermediate host bus 119 to which CPUs 111 are connected, a host bus bridge 114, and a peripheral bus 118, such as a PCI bus, to which NIC 112 is connected. In one embodiment, NIC 112 includes its own special purpose or embedded processor to enable NIC 112 to perform some local processing of network frames.

For embodiments of service device 110 that include multiple processors 111 or that include a NIC 112 with sufficient processing capability, server device 110 possesses the ability to simultaneously perform at least two processes. Generally speaking, the invention contemplates using the processing capability of multi-processors servers and/or servers with network cards containing hardware-assist to reduce latency in responding to web based or other network requests.

Referring now to FIG 3, an illustration of the operation of server device 110 in responding to a client initiated request for information is depicted. In the depicted embodiment, a frame 200 representing a client initiated request, or a portion thereof, is indicated as arriving at

server device 110 presumably via NIC 112. NIC 112 preferably includes sufficient processing capability and firmware to perform local processing of frames 200 received from network 104.

Upon receipt by server device 110 and NIC 112, frame 200 is initially provided to a frame parser 306. Frame parser 306 may comprise firmware, software, or hardware implemented on NIC 112. Alternatively, frame parser 306 may comprise software executed by one of the multiple processors 111 in a multiprocessor implementation of server 110. In either case, frame parser 306 is responsible for defining a first portion of frame 200, referred to herein as the network portion 301, and a second portion of frame 200 referred to herein as the data portion 302. As their names suggest, network portion 301 of frame 200 includes the network relevant information such as the network layer header 208 and the transport layer header 206 while the data portion 302 of frame 200 includes the request relevant information including the application layer header 204 and any data 202. In addition, network portion 301 may include part or all of data portion 302 or even the entire frame 200 in an embodiment where the error checking mechanism relies upon the frame as whole.

In a TCP/IP network environment in which server 110 is handling web-base requests, the network portion 301 of frame 200 includes the TCP and IP headers. The IP header information is used to determine, among other things, whether frame 200 arrived at the correct destination (i.e., was the frame properly routed over network 104). The TCP header information is used to ensure the integrity and reliability of the frame. Although processing these header can require an appreciable amount of time, such as the time required to perform a full ECC check of frame 200, frame parser 306 can delineate relatively quickly the network portion 301 from the data portion 302.

Network portion 301 of frame 200 is then provided to a frame verifier represented in FIG 3 by reference numeral 304. Frame verification 304 may exist as software or firmware on NIC 112. The frame verification 304 may include processing the network layer header 208 and a transport layer header 206 of a frame 200 to determine if the frame 200 is in the proper location and whether it contains any errors.

Simultaneously with the processing of network portion 301 of frame 200 in the frame verification unit 304, the data portion 302 of frame 200 is being processed by the server application. In an embodiment in which server device 110 is servicing web-based requests, server application 303 may represent a web server that interprets HTTP formatted Hypertext

Markup Language (HTML) code. Typically, server application 303 will be responsible for retrieving data from a storage facility such as a hard disk in response to the client initiated request. Typically, the retrieval of this information is unaffected by the network portion information and can therefore proceed in parallel with the processing of network information

5 301. Ultimately, server application 303 produces requested data 306 responsive to the client request and forwards requested data 306 to a frame formatter 308.

In the depicted embodiment, the frame verification unit 304 acts as an enable for frame formatter 308. More specifically, if the frame verification unit 304 determines that the frame 200 received by server device 110 is reliable, the verification unit permits frame formatter 308 to

10 initiate the process of formatting the requested information suitable for transmission across the network. If frame verification unit 304 detects an error condition of some sort, further processing of the request data is terminated.

If the received frame 200 is verified as reliable, frame formatter 308 is responsible for formatting the requested data 308 into protocol compliant frames suitable for transmission across network 104 back to the requesting client 102. By optimistically assuming that the received frame 200 is reliable (contains no errors and is appropriately routed) the invention reduces latency by trading incurring the relatively small delay required to parse the frame in exchange for the ability to process the network information and the request information simultaneously. Since the majority of frames are presumed to be reliable even in an Internet application, the tradeoff will effectively reduce latency and improve performance.

15 20

It will be apparent to those skilled in the art having the benefit of this disclosure that the present invention contemplates a method and system for responding to client initiated request with reduced latency. It is understood that the form of the invention shown and described in the detailed description and the drawings are to be taken merely as presently preferred examples. It

25 is intended that the following claims be interpreted broadly to embrace all the variations of the preferred embodiments disclosed